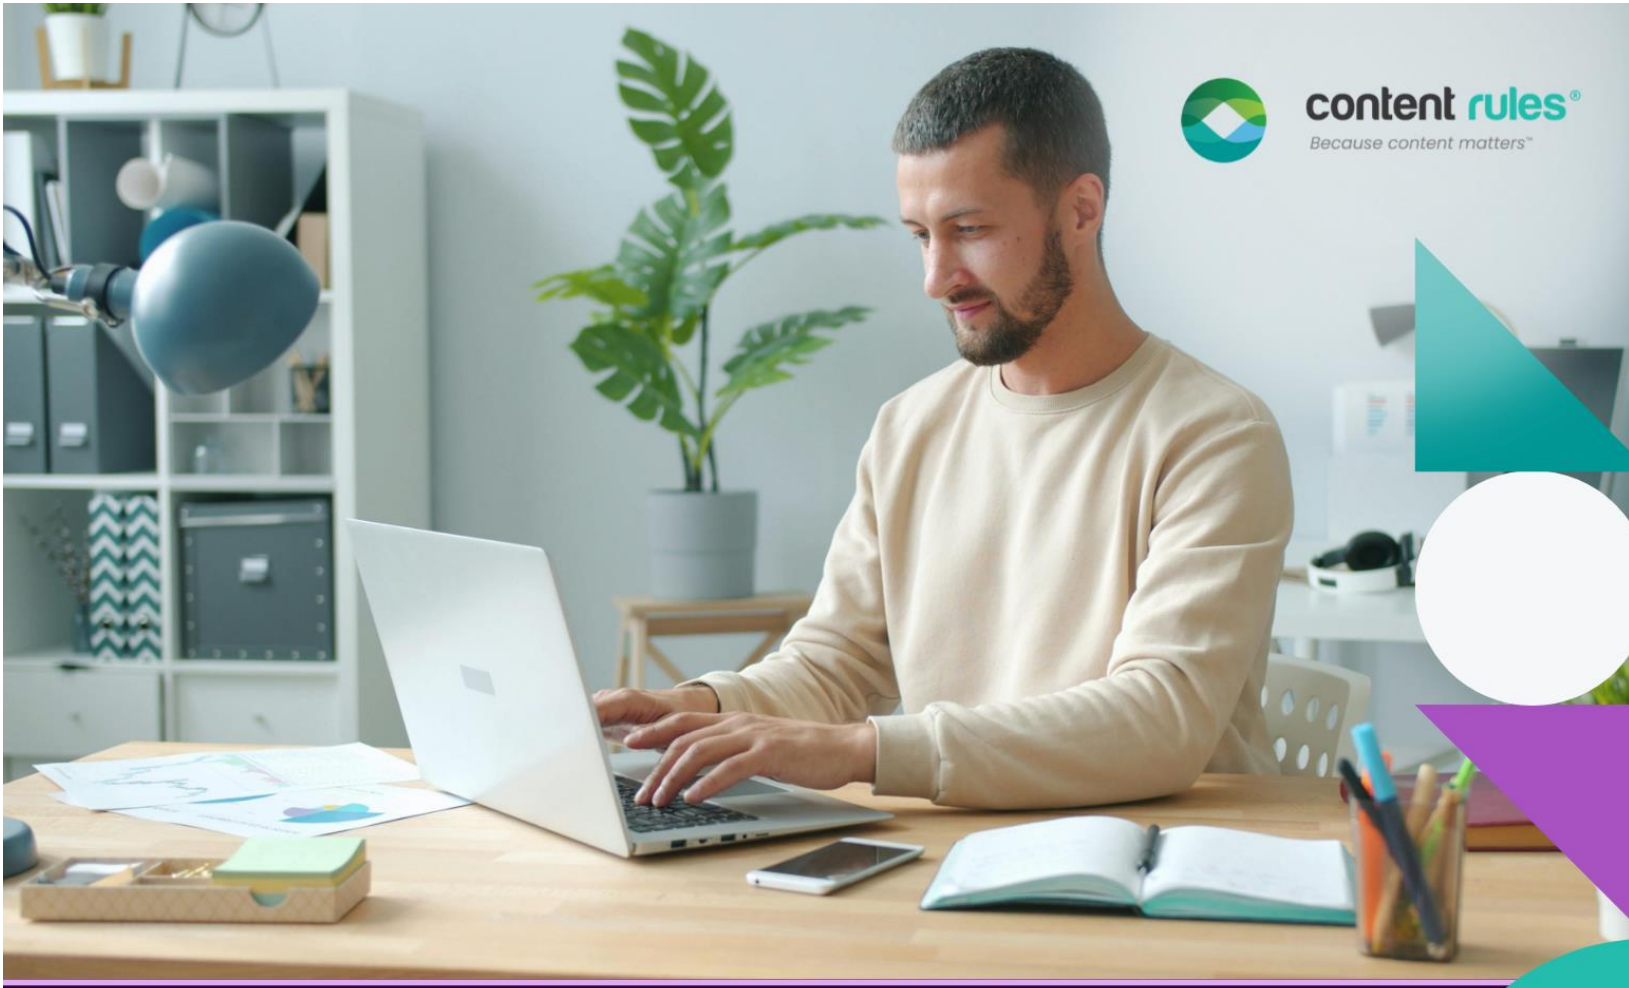


content rules®
Because content matters™

WHITE PAPER

From Documents to Knowledge

How a knowledge first mental model enables digitalization, automation, and AI

© Content Rules, Inc. All Rights Reserved.

Table of Contents

Knowledge First is the missing link 3

What is a mental model? 4

What is the Knowledge First mental model? 5

Knowledge is the business asset 6

Capture, curate, and communicate 7

- 1. Capture knowledge 7
- 2. Curate knowledge 8
- 3. Communicate knowledge 9

Locking knowledge in documents is bad for business and for patients 10

- Risks of document-based systems 10
- Documents limit AI reliability 12
- Human-in-the-loop needs a digital approach 13

Digitalization requires a new mental model 15

About the authors 16

References 17

About Content Rules 18

© Content Rules, Inc. All Rights Reserved.

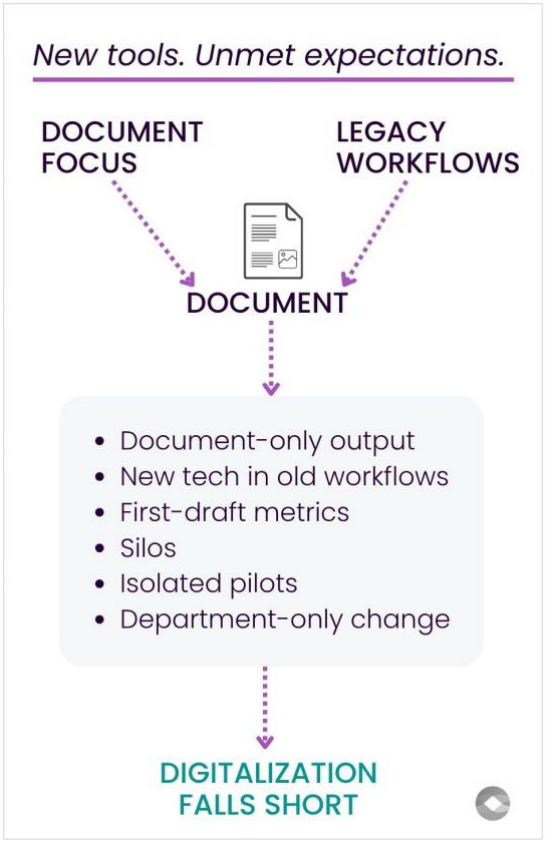
info@contentrules.com | contentrules.com

Knowledge First is the missing link

Organizations have tried several ways to leverage technology to accelerate content workflows. Unfortunately, few have realized the benefits they expected in shorter timelines, greater efficiency, or reduced costs.

Here are some common reasons why their expectations are not met:

- Focus on generating documents instead of building knowledge assets
- Replicate legacy workflows in new technology
- Measure "time to first draft" instead of assessing the end-to-end knowledge lifecycle
- Allow legacy silos to block technical and data enhancements
- Operate in pilots without addressing underlying issues that prevent scalable processes
- Treat digitalization as a departmental initiative rather than an organization-wide transformation



The common thread that ties these disappointed expectations together is not the technology. It's not even the process.

The number one reason transformation fails is that there has been no lasting change to the mental model behind how content is created, managed, and communicated. These organizations continue to focus on documents. Instead, they need to focus on knowledge.

© Content Rules, Inc. All Rights Reserved.
info@contentrules.com | contentrules.com

What is a mental model?

The exact definition of *mental model* varies depending upon which discipline you approach it from. In this paper, *mental model* is defined as "an interrelated set of beliefs that shapes a person's expectations for the future and how they understand the way the world works."[1]

Without a new mental model of how knowledge is captured, curated, and communicated, most people simply cannot make digital transformation successful. They continue to think in documents. They hold on to legacy processes instead of embracing more effective ways for their expertise to contribute to the business.

When that happens, the organization cannot realize the full potential of technology to shorten timelines, improve efficiency, and reduce risk. The return on investment (ROI) looks weak and the technology takes the blame. The digital transformation falls short of its potential or even fails altogether.

This paper describes the mental model that makes the difference between successful digital transformations that deliver business value and other attempts that do not fully pay off.

We call this mental model Knowledge First.

© Content Rules, Inc. All Rights Reserved.
info@contentrules.com | contentrules.com

What is the Knowledge First mental model?

Knowledge First is a visionary approach to capturing, optimizing, and delivering knowledge for human use and machine processing. The Knowledge First mental model changes the picture from "we must produce documents" to "we must capture knowledge." This approach enables large organizations to achieve business objectives more efficiently and with less risk.

In the Knowledge First mindset, writers and subject-matter experts (SMEs) focus on creating *knowledge assets* that support multiple functions within the business. A knowledge asset is a unit of information that is augmented for machine processing so that automation can be implemented to speed up processes and to reduce errors.

Knowledge assets typically comprise data, evidence, and meaning that support the knowledge lifecycle across the organization. Rather than being limited to a single document or use case, this knowledge becomes available to people throughout the organization to make decisions, evaluate outcomes, and share results.

When that knowledge is locked away in a document that is stored with hundreds (or thousands) of other documents or saved in a team's network drive, it becomes difficult or impossible to find and use that knowledge. As a result, people can make decisions and take actions based on incomplete or inaccurate information. These decisions incur risk, cost, and time that could have been avoided if the knowledge had been more readily available.

This risk increases when we bring AI into the equation.

Just about everyone knows that the quality of AI results depends almost entirely on the quality of the source information. But not everyone realizes that documents are not the optimal source format for AI to access knowledge.

© Content Rules, Inc. All Rights Reserved.
info@contentrules.com | contentrules.com

Knowledge is the business asset

> *The crux of the Knowledge First mental model is that knowledge is the primary business asset. The document is not.*

The Knowledge First mental model helps people understand that their knowledge is much more valuable to the business when it is optimized for systems to find, retrieve, and use. Their effort is rewarded when the knowledge assets they contribute are referenced by parallel and downstream systems to accelerate or de-risk business processes across the organization.



The concept of knowledge assets rather than documents is becoming more and more critical as organizations deploy agentic AI workflows. These agents are intended to take actions based on the knowledge that has been captured. Data alone is not enough to enable trustworthy agentic AI workflows; structured knowledge is key to building autonomous agentic AI at scale.[2]

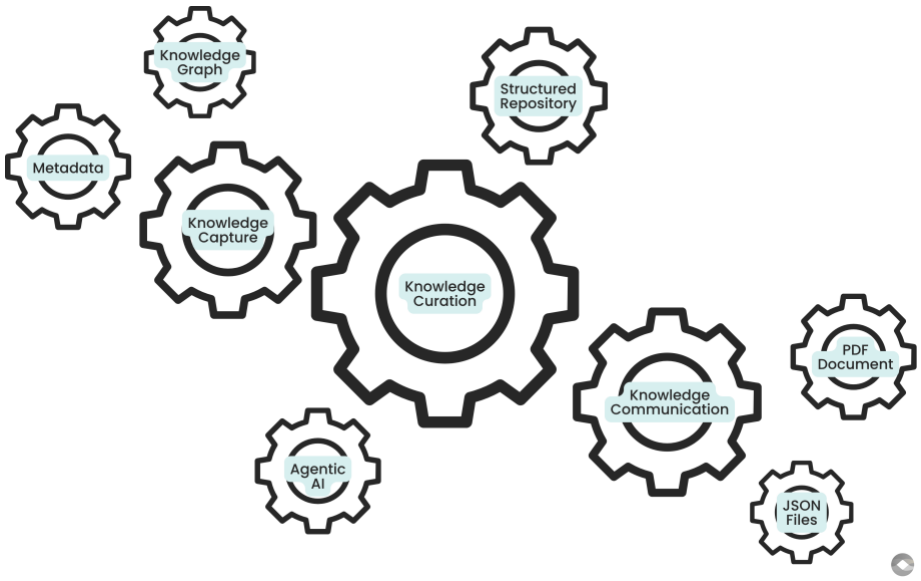
McKinsey found that only seven percent of business leaders surveyed in 2025 reported that AI was fully scaled within their organizations.[3] The report says that scaling AI requires more than adopting new technology. According to McKinsey analysts, businesses need to redesign workflows around AI capabilities and build platforms that allow them to operate at scale.

© Content Rules, Inc. All Rights Reserved.
info@contentrules.com | contentrules.com

The enterprise goal of leveraging AI to interpret information and take actions may seem removed from the daily work of any given individual involved with knowledge development. And yet, the success of any AI implementation beyond a carefully constructed proof of concept depends on making the radical shift in how the source knowledge is identified, created, and managed over the long term.

Capture, curate, and communicate

The Knowledge First mental model divides the digitalization of enterprise knowledge into three categories: capture, curate, and communicate.



1. Capture knowledge

The Knowledge First mental model envisions capturing knowledge at the genesis point of the information. The knowledge is then stored and classified to enable reuse, transformation, and referencing any time the information is needed.

© Content Rules, Inc. All Rights Reserved.

That information must be findable and usable for a very long time. Organizations continually create knowledge across products, services, customers, operations, and other areas of the business. These knowledge assets can include data, evidence, analysis, proprietary intellectual property, and other information that must remain available as the business evolves.

Once identified, these knowledge assets need to be classified and labeled so they can be identified, discovered, and maintained. Documents simply do not allow teams to label knowledge assets at the level of granularity required to enable fast, accurate search results or support AI to produce reliable results.

In the Knowledge First model, this capture and identification of the knowledge asset is a critical factor for success.

2. Curate knowledge

The Knowledge First mental model envisions knowledge curation as the ongoing process of keeping digital knowledge assets current, correct, and connected throughout their lifecycle.

It is (relatively) easy to curate the knowledge one time. The real challenge is in the ongoing maintenance of that knowledge.

Most digital transformation initiatives have focused first on automating the creation of new documents. The core project team is tasked with conducting proof of concepts or pilot programs for a specific objective or use case. Maintaining and updating knowledge is often a secondary concern, if considered at all.

When the use case is to update or expand the existing knowledge within the digital process, the question of knowledge curation can no longer be put off. That's when the transformation program is often abandoned as the full end-to-end knowledge curation process has not been addressed.

© Content Rules, Inc. All Rights Reserved.

The Knowledge First mental model emphasizes curation as an equal priority to capturing and communicating knowledge. Curation is the hub of the entire mental model.

3. Communicate knowledge

The Knowledge First mental model envisions knowledge communication as an end result of capturing and curating knowledge.

The same digital knowledge assets can be assembled and transformed into any number of communication formats.

A document is one traditional format for communicating knowledge. Other formats can include **web content, structured data, application interfaces, knowledge bases, AI-generated responses, and information** exchanged directly between systems.

These knowledge communication formats happen downstream from knowledge creation. Many teams are not aware of how their knowledge is used or that their knowledge contributes to these processes at all.

Today, as AI becomes more prevalent in enterprise workflows, the availability of knowledge moves to the forefront of almost every business endeavor. Business processes that used to happen "downstream" now happen in parallel or even move upstream, occurring before a document is produced. This shift in workflow goes more smoothly when people make a shift in how they envision the end-to-end knowledge lifecycle.

The Knowledge First mental model ensures that the people who capture or curate knowledge do not need to know all the formats that will be used to communicate knowledge. Rather, the mental model requires people to think of the communication format as an "output" of knowledge that is separate from the "input" of knowledge. We capture and curate knowledge through the inputs; we communicate through the outputs.

© Content Rules, Inc. All Rights Reserved.

This separation of inputs and outputs is a relatively new concept to most organizations. Yet it is one of the key distinctions between legacy workflows (which are based on how paper works) and digital workflows (which are based on how computers work).

Locking knowledge in documents is bad for business

Knowledge is a critical enterprise asset for organizations of every size.



Knowledge needs to be findable in order to be valuable. And yet organizations regularly lock their knowledge away in the least accessible format for both humans and machines: the document.

Risks of document-based systems

Document-based systems have long been impediments to communicating knowledge. This fact is not new.

Many organizations now have decades of experience capturing knowledge in digital assets that can be searched, assembled, and delivered in any number of communication formats.

Dell, Inc.'s technical communications team had already been creating structured knowledge assets for more than 10 years when agentic AI became a business reality.[4] Those knowledge assets were captured, stored, classified, and curated to enable about 60% content reuse across millions of HTML and

© Content Rules, Inc. All Rights Reserved.

PDF pages. In 2026, these assets became the foundation of the company's "Next Best Action" AI agent. This solution greatly speeds up customer support resolution and is credited **with an estimated $240 million in support cost deflection.**[4] Their workflows reflect the Knowledge First mental model and quickly adapt to support AI solutions.

Yet many organizations still struggle to make a similar shift in how knowledge is captured, managed, and delivered.

As a result, organizations cannot:

- Amass knowledge into an accessible place while also preventing redundancy, contradictions, and inconsistencies
- Keep information current and accurate without extensive manual work that is often hidden and not documented
- Leverage automation to provide the right content in the right format
- Provide enough context for AI to find the exact right information in a document library riddled with redundancy and potential contradictions

Unfortunately, some of those organizations have attempted to address these risks by implementing AI without taking a critical look at the knowledge foundation.

© Content Rules, Inc. All Rights Reserved.

Documents limit AI reliability

We can all see the potential of AI and content automation to accelerate how organizations create, access, and use knowledge. And we have seen over and over again that new tools alone do not deliver as much acceleration as desired.

> *AI cannot reach its full potential when knowledge is locked away in unstructured documents.*

Organizations are increasingly looking to AI as a way to speed up document-based processes. However, the risks inherent in document-based processes increase, not decrease, when you add AI without also changing how you capture and curate knowledge. In fact, according to Gartner, more than 40% of agentic AI projects are likely to be canceled by the end of 2027.[5]

Nobody knows just how many errors, gaps, contradictions, and outdated information exist in the vast document repositories that represent an organization's accumulated knowledge. That was accepted as an inevitable part of the document-centric process when the risks lay in the chance that (1) a person encountered one of these issues and (2) it caused enough of a problem to require escalation.



In test after test, **structured knowledge assets** enriched with context and explicitly defined relationships **outperform static documents** in any AI-enabled processing.

30%-60%
Faster
Structured graph queries provide 30% to 60% faster initial generation than ad-hoc chunk retrieval.[6]

Context
Matters
RAG retrieval suffers when source content lacks context, metadata, and explicit relationships.[7]

Structured Data
Required
It is impossible to automate classification tasks reliably without structured data.[2]

Less
Training Data
AI models trained on well-labeled and semantically defined content perform better with less training data.[8]

When AI is incorporated into document-based processes, the errors, gaps, and outdated information within thousands of documents become obvious on a scale that wasn't conceivable with search results and human assessment. These issues can affect decisions, customer experiences, operations, and other business processes that depend on accurate, trustworthy knowledge.

© Content Rules, Inc. All Rights Reserved.

With AI, there is no hiding those deficiencies. Can you detect them before they propagate and escape to the outside world? What are the repercussions for customers and other stakeholders? To enterprise reputation? To the bottom line?

Organizations expect to use AI to transform source knowledge assets into additional assets for use across reports, customer-facing content, digital experiences, and other outputs. The expectation is that AI will help teams create and communicate knowledge more efficiently and at scale.

This expectation can only be realized when AI applications and agents have access to curated, structured knowledge assets.

Human-in-the-loop needs a digital approach

One common approach to digitalization is to incorporate "human-in-the-loop" (HITL) review as the safety net for AI-generated documents.

The reliance on HITL reflects the reality that we cannot trust AI to do the job correctly when we use AI to generate documents from unstructured knowledge sources. However, Dell's content engineering approach suggests that when humans focus on capturing and curating structured knowledge assets, then AI can produce much better results.[4]

The combination of AI + HITL is supposed to cut weeks and months out of the document development process. It can, but not if you just use documents to create more documents. The process of reviewing information still requires review of source documentation to ensure that the analysis and conclusions are consistent and complete. People must also perform close review of the language and details so that the output presents a clear, accurate, and persistent view of the underlying information and conclusions.

© Content Rules, Inc. All Rights Reserved.

A digital approach that puts structured knowledge assets at the center is key to achieving process acceleration and content quality at scale.

> *Organizations that have achieved significant acceleration in content development did not achieve these results simply from adding AI to document-based processes. Rather, they changed their fundamental approach from documents to structured knowledge assets. They adopted new practices such as lean authoring, structured content authoring, structured content and data management, and lean templates.*



Structured knowledge creates measurable business value across industries.

© Content Rules, Inc. All Rights Reserved.
info@contentrules.com | contentrules.com

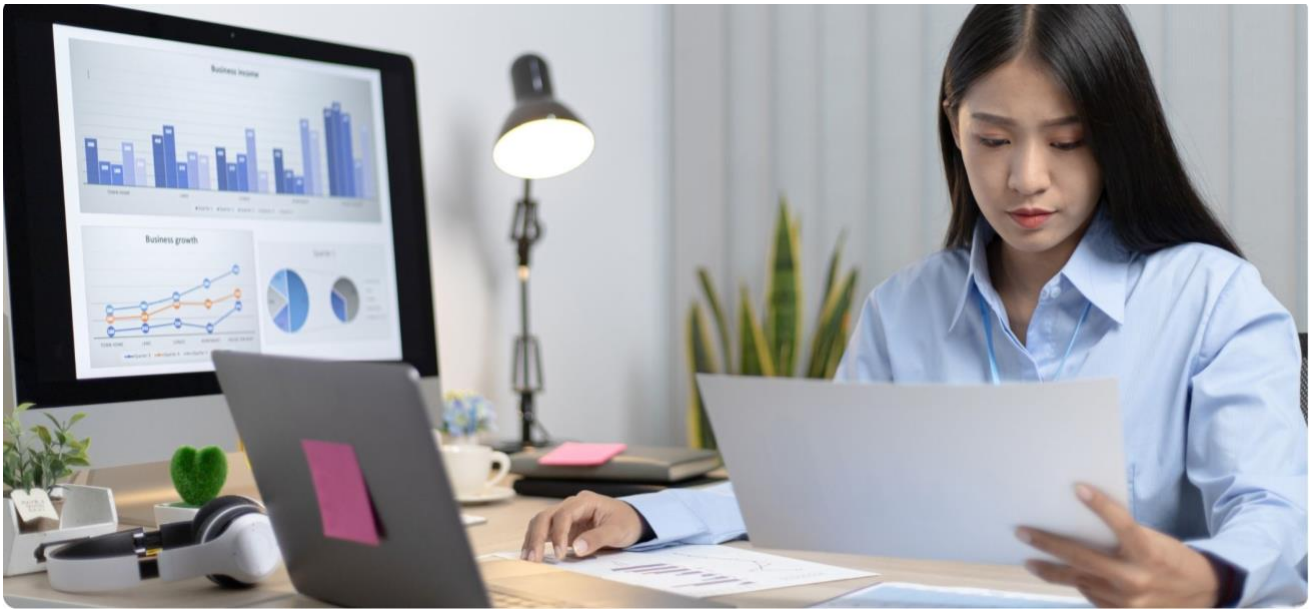
Digitalization requires a new mental model

Adding new technologies into existing document-first processes or undertaking single-purpose pilots that do not scale will not equate to successful digitalization. Realizing the full ROI requires a significant change in how people work. For people to accept these changes, they must develop a new understanding of digitalization.

To be successful at scale, organizations must fundamentally change their current approach to digitalization. It is evident that the current approaches are not providing the desired results.

The fundamental change required is to change the way we think about knowledge itself. Knowledge First is the mental model that makes that shift possible.

The Knowledge First mental model helps the people closest to the business-critical knowledge develop a common vision. This unifying vision then enables them to adopt new technologies, processes, and workflows that empower successful digitalization.



© Content Rules, Inc. All Rights Reserved.
info@contentrules.com | contentrules.com

About the authors



Kathleen Vieson

Kathleen Vieson, PharmD, CMPP, is a life sciences content strategy and digital transformation consultant. She helps pharmaceutical companies and content technology vendors transform clinical and drug information through structured content, governance, technology implementation, and AI readiness. A clinical pharmacist and oncology specialist by training, Kathleen brings deep expertise across healthcare information, content operations, and technology. She has led large cross-functional content teams, developed structured content and reuse strategies, and built content and data management systems that support scalable digital delivery. Today, she advises global pharma organizations on moving from document-centric processes to structured, reusable knowledge ecosystems.



Regina Lynn Preciado

Regina Lynn Preciado is the VP of Content Strategy Solutions at Content Rules, Inc. She leads content strategy teams to help pharma organizations optimize content workflows and maximize the business value of their investment in structured content ecosystems. Regina has helped clinical, non-clinical, labeling, quality, medical device, and marketing teams mitigate risk and enable sustainable content practices that scale. Regina is an industry expert with more than 25 years of experience in structured content authoring, component content management, AI readiness, and content reuse and automation. She lives a dogspotting lifestyle.

© Content Rules, Inc. All Rights Reserved.
info@contentrules.com | contentrules.com

References

[1] Holtrop JS, Scherer LD, Matlock DD, Glasgow RE, Green LA. The Importance of Mental Models in Implementation Science. *Front Public Health*. 2021 Jul 6;9:680316. doi: 10.3389/fpubh.2021.680316. PMID: 34295871; PMCID: PMC8290163. https://pmc.ncbi.nlm.nih.gov/articles/PMC8290163/#s2

[2] Talisman, J. (2025, June 30). Why AI isn't autonomous (yet), Part III: What agentic AI needs. Intentional Arrangement. *Substack*. https://jessicatalisman.substack.com/p/why-ai-isnt-autonomous-yet-ea2

[3] McKinsey & Company. (2025, November). AI at work but not at scale. *McKinse.y* https://www.mckinsey.com/featured-insights/week-in-charts/ai-at-work-but-not-at-scale

[4] Eric Brehmer (2026). The impact of agentic AI on technical documentation: How Dell is responding. Tridion Live Sessions, https://www.rws.com/content-management/tridion/resources/tridion-live-sessions-2026/

[5] Gartner, Inc. (2025, June 25). Gartner predicts over 40% of agentic AI projects will be canceled by end of 2027. *Press release*. https://www.gartner.com/en/newsroom/press-releases/2025-06-25-gartner-predicts-over-40-percent-of-agentic-ai-projects-will-be-canceled-by-end-of-2027

[6] Lantosca, M. (2025, October 26). Using knowledge graphs to accelerate and standardize AI-generated technical documentation. *Medium*. https://medium.com/@nc_mike/using-knowledge-graphs-to-accelerate-and-standardize-ai-generated-technical-documentation-for-03f07b8df585

[7] Choubey, P. K., Peng, X., Bhagavath, S., Huang, K.-H., Xiong, C., & Wu, C.-S. (2025). Benchmarking deep search over heterogeneous enterprise data. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing: Industry Track* (pp. 501–517). Association for Computational Linguistics. https://aclanthology.org/2025.emnlp-industry.34/

[8] Geary, L., Nielsen, I.Ø., Charron, F., Carter, J. J., Martinsson, L., Rowe, E., Goese, A., Chaturvedi, A., Gidh-Jain, M., Hreha-Phillips, S., & April, J. (2026). Optimizing clinical study report development through lean authoring: A framework for meeting regulatory expectations and preparing for automation. *Therapeutic Innovation & Regulatory Science*. https://link.springer.com/article/10.1007/s43441-026-00980-6

[9] Merck & Co. (2021). Clinical study report optimization and reduction of document redundancy. Journal of Pharmaceutical Sciences. https://doi.org/10.1016/j.xphs.2021.09.046

[10] RWS. (n.d.). HORIBA Medical scales its documentation processes. https://www.rws.com/content-management/tridion/resources/horiba-medical-case-study/

[11] MadCap Software. (n.d.). Avalara, a leading tax compliance services company, leveraged IXIA CCMS to unify their internal and external content on a single platform. https://www.madcapsoftware.com/case-study/avalara-uses-madcap-ixia-ccms/

[12] Componize. (n.d.). Streamlining service bulletins at Toyota Motor Europe. https://componize.com/case-studies/toyota-motor-europe/

© Content Rules, Inc. All Rights Reserved.

info@contentrules.com | contentrules.com

About Content Rules

Content Rules, Inc. helps pharma organizations accelerate adoption and maximize the return on their investment in digitalization. The world's largest and most innovative companies trust us to develop their enterprise content strategies, transform their content ecosystems, optimize content for a worldwide audience, and develop effective content that gets results.

Since our founding in 1994, we have refined our approach over decades of working with organizations facing challenges just like yours.

We provide:

- Knowledge First training, coaching, and mentoring
- Content strategy and AI readiness assessments
- Content architecture and reuse strategy design
- Writer and reviewer enablement
- Governance and workflow design
- Pilot and implementation support
- Long-term adoption guidance

Contact us at info@contentrules.com for a free consultation.

© Content Rules, Inc. All Rights Reserved.